

RECEIVED
CENTRAL FAX CENTER

NOV 17 2006

REMARKS

In the Office Action, the Examiner rejected the claims under 35 USC §103. The rejections are fully traversed below. The claims have been amended correct various typographical errors. Claims 1, 3, and 5-61 remain pending.

Reconsideration of the application is respectfully requested based on the following remarks.

REJECTION OF CLAIMS UNDER 35 USC §103(a)

In the Office Action, the Examiner has rejected claims 1, 12, 14, 15, and 18-53 under 35 USC §103(a) as being unpatentable over Blumenau et al, U.S. Patent No. 6,260,120, ('Blumenau' hereinafter) in view of Oberman et al, U.S. Pub. No. 2003/0026267 ('Oberman' hereinafter). The Examiner has also rejected claims 29-48, 50-53, and 59-61 under 35 USC §103(a) as being unpatentable over Blumenau and Oberman. This rejection is fully traversed below.

Various embodiments of the invention support the virtualization of storage in a storage area network. This is accomplished through the use of one or more network devices capable of being placed in a data path between the hosts and the storage devices. As a result, neither the storage devices nor the hosts require additional software or hardware to support storage virtualization. Moreover, multiple network devices may simultaneously manage the virtualization of heterogeneous storage devices.

The pending claims implement storage virtualization, as claimed, on a per-port basis. In other words, selected ports of one or more network devices may implement virtualization functionality in hardware and/or software. Any number of ports on a switch can manage

virtualization of its own traffic. This allows a network's virtualization capacity to scale with the number of ports.

Blumenau relates to storage mapping and partitioning among multiple host processors in the presence of login state changes and host controller replacement. See title. A storage controller is programmed to define a respective specification for each host processor of a respective subset of the data storage to which access by the host processor is restricted, and each specification is associated with a host identifier stored in the memory. When the storage controller receives a data access request from a host processor, it decodes a host identifier from the data access request, and searches the memory for a host identifier matching the host identifier decoded from the request. Upon finding a match, the respective specification of the respective subset for the host processor is accessed to determine whether or not storage specified by the storage access request is contained in the respective subset. If so, then storage access can continue, and otherwise, storage access is denied. Preferably, the host identifier decoded from the request is a temporary address assigned by the network, and also stored in the memory in association with each respective specification is a relatively permanent identifier for the host processor. See Abstract.

As shown in Fig. 22 of Blumenau and described in col. 26, lines 1-24, a cached storage subsystem 250 implementing virtual ports 268 includes two port adapters 260 and 261, each having two physical ports. The port adapters are programmed to provide respective virtual switches linking their physical ports to a set of virtual ports. As shown in Fig. 7, the cached storage subsystem 20 appears to include a single network device having a cache memory 32 accessible by both port adapters. This is further emphasized in col. 9, lines 25-29, stating "in a preferred form of construction, the cache memory 32 is composed of dynamic RAM memory cards mounted in a card-cage or main-frame, and the port adapters and storage adapters are programmed micro-processor cards that are also mounted in the card-cage or main-frame."

As set forth above, Blumenau neither discloses nor suggests implementing storage virtualization, as claimed, on a network device. While Blumenau does disclose the mapping of LUNs to logical volume numbers (see col. 22, lines 60-67), Blumenau fails to disclose the claimed inventions as implemented in a network device, wherein the network device is a switch, router, iSCSI gateway, or other network node configured to perform a switching function.

The Examiner admits that Blumenau "fails to teach wherein (b), (c), and (d) are performed by logic dedicated to and implemented by said port of the network device." The Examiner seeks to cure the deficiencies of Blumenau with Oberman. Oberman relates to virtual channels in a network switch. See title. In the rejection, the Examiner cites paragraph [0148] of Oberman, which states:

In one embodiment that supports virtual channels, there are K virtual channels available per port. In one embodiment, K=8. In one embodiment, the slice keeps track of the resources (e.g. packets and clusters) in use by every virtual channel on every port of the slice. Resources may be allocated to every virtual channel on every port using programmable registers. When a packet comes into the slice, the packet admission logic may check to see whether the resources in use by the virtual channel to which the packet belongs are less than the allocated limits. Both packet descriptors and clusters may be checked. If either of these checks fail, the packet may be dropped. The packet may also be dropped if either a packet descriptor or a cluster is not available. Early Forwarding may be allowed for virtual channel packets if the total available resources for the virtual channel on the port are greater than or equal to a programmable value (e.g. ClmMinFreeClustersEarlyFwdVC_x_Port_y, where x is the virtual channel number and y is the port number). Every virtual channel on every port may have a different packet size negotiated with its link partner. Hence, in one embodiment, there are individual registers for each virtual channel on each port.

It is important to note that the cited portion of Oberman merely relates to "resource tracking." In other words, Oberman merely discloses the allocation of resources to each virtual channel on every port using programmable registers. While Oberman discloses individual registers for each virtual channel on each port, Oberman does not disclose performing "(b) determining that the frame or packet pertains to access of a virtual storage location of a virtual storage unit representing one or more physical storage locations on one or more physical storage units of the storage area network; (c) obtaining a virtual-physical mapping between the one or more physical storage locations and the virtual storage location; and (d) sending a new or modified frame or packet to an initiator or a target specified by the virtual-physical mapping" where (b), (c) and (d) are performed by logic dedicated to and implemented by a port of a network device. Accordingly, the combination of the cited

references would fail to achieve the desired result.

Paragraph [0014] of Oberman details the problem in the art that Oberman attempts to solve:

Because of unreliable communication (bit errors may eventually occur), packets on a link may be corrupted and as such "lost". If packets that include credit information become corrupted, credits may be lost, potentially resulting in a deterioration of transmission rate. Eventually, if all the credits are lost, transmission of packets over the link will stop. Fiber Channel Arbitrated Loop (FC-AL) avoids this problem by refreshing the credits to the initial number every time a port is opened. Fibre channel point-to-point connection has no specific mechanism to do this, other than the credits being refreshed when the port gets into an "error" state. In order to avoid this problem, it is desirable to provide a credit synchronization procedure for network switches implementing credit-based flow control for storage packet flows on links between the switches.

Thus, Oberman merely attempts to avoid the loss of credits in a credit-based system.

The problem being solved by Oberman is entirely different from that of the claimed invention. Nothing in the cited art, separately or in combination, discloses or suggests implementing virtualization functionality on selected ports of one or more network devices. Stated another way, the cited art, separately or in combination, fails to disclose or suggest enabling any number of ports on a switch to manage virtualization of its own traffic. Moreover, the combination of the cited references would fail to achieve the advantages of the pending claims, which enable a network's virtualization capacity to scale with the number of ports in a network. In view of the above, Applicant respectfully requests that the Examiner withdraw the rejection of claims 1, 12, 14, 15, and 18-53 under 35 USC 103.

In the Office Action, the Examiner has rejected claims 3, 5-11, 13, 16, and 17 under 35 USC §103(a) as being unpatentable over Blumenau and Oberman, and further in view of

Lo et al, U.S. Pub. No. 2002/0103943 ('Lo' hereinafter). This rejection is fully traversed below.

The Examiner admits that Blumenau and Oberman fail to teach a network device, where the virtual storage unit comprises a VLUN or other virtual representation of storage on a storage area network. The Examiner seeks to cure the deficiencies of Blumenau and Oberman with Lo.

Lo discloses a distributed storage management platform architecture. See title. Lo does disclose the concept of storage virtualization. See paragraphs 0037, 0239. The behavior disclosed in Lo is storage network router-based, rather than being housed in the hosts, or in the storage arrays/subsystems. See paragraphs 0247, 0048-0049.

In no manner does Lo disclose or suggest performing mapping functionality within a system implementing virtualization of storage on a per-port basis. In fact, Lo discloses that functionality other than conversion between frame formats be performed by other entities. For instance, col. 19, lines 38-52 disclose that storage commands are directed to a storage traffic engine, which handles storage commands. Moreover, storage traffic is sent to the storage traffic block. In other words, these commands and storage traffic are not handled by a port. Accordingly, Lo teaches away from performing or obtaining a virtual-physical mapping in a system implementing virtualization on a per-port basis.

It is important to note that the combination of the cited references would fail to achieve the desired result. Even if conversion between formats is supported by the port interfaces, any mapping performed within a system implementing storage virtualization would not be performed on a per-port basis. Accordingly, the combination of the cited references would fail to operate as claimed.

As set forth above, the claimed invention enables any number of ports of one or more network devices within the network to manage virtualization of its own traffic. As a result, the network's virtualization capacity scales with the number of ports. Since Lo fails to disclose or suggest performing virtual-physical mapping on a per-port basis, the network's virtualization capacity cannot match the capacity of the claimed invention. Accordingly, the claimed invention yields advantages that are not possible with the cited art.

While Lo discloses the concept of storage virtualization, Lo fails to cure the deficiencies of the primary references, as set forth above. Accordingly, Applicant respectfully

requests that the Examiner withdraw the rejection of claims 3, 5-11, 13, 16, and 17 under 35 USC 103.

In the Office Action, the Examiner has rejected claim 49 under 35 USC §103(a) as being unpatentable over Blumenau and Latif et al, U.S. Patent No. 6,400,730 ('Latif' hereinafter), and further in view of Lo. This rejection is fully traversed below.

As disclosed in the abstract of Latif, "the port interfaces provide the conversion from the input frame format to an internal frame format, which can be routed within the apparatus." Examples of such conversion set forth in Latif include data switching between SCSI and IP, Fibre Channel and IP, and SCSI and Fibre Channel. Moreover, Lo teaches a network device, where the type of traffic is iSCSI. (See paragraph 0128).

However, as set forth above, Lo teaches away from obtaining or performing a virtual-physical mapping on a per-port basis. Accordingly, Applicant respectfully submits that claim 49 is patentable over the cited references.

In the Office Action, the Examiner has rejected claims 18 and 19 under 35 USC §103(a) as being unpatentable over Blumenau and Oberman, and further in view of Latif et al, U.S. Patent No. 6,400,730 ('Latif' hereinafter). This rejection is fully traversed below.

The Examiner admits that Blumenau "fails to teach wherein (b), (c), and (d) are performed by logic dedicated to and implemented by said port of the network device." The Examiner cites col. 18, lines 8-42 of Latif, stating that Latif teaches a network device, wherein (b), (c), and (d) are performed by a processor dedicated to only said port of the network device.

Col. 18, lines 8-42 of Latif indicates that the routing logic includes logic blocks that are dependent on the port type and other blocks that are common to different port types. Routing block is described as determining where a frame is routed based upon addressing information within the data frame. Latif further indicates that address resolution logic can be

shared by two different port interfaces. In addition, Latif indicates that routing logic block 350 is shared by two port interfaces. Moreover, Latif further indicates that additional logic can also be shared. Thus, Latif does not require that the routing logic be dedicated to a single port, but rather teaches that routing logic is to be shared among multiple ports. As such, Latif teaches away from the claimed invention. It is also important to note that the cited portion of Latif does not explicitly require that the routing logic include claimed steps (b), (c), and (d). Accordingly, Applicant respectfully submits that claims 18-19 are patentable over the cited art.

The Examiner has also rejected claims 21-28 under 35 USC §103(a) as being unpatentable over Blumenau, Oberman, Latif, and further in view of Brewer et al, U.S. Patent No. 6,876,656, ('Brewer' hereinafter). This rejection is fully traversed below.

As claimed in claims 21-28, a port may submit a lock request to another "master" port that manages all lock requests. This master port may also notify the requesting port when a lock request has been granted. The master port may similarly process "lock release" requests.

With respect to claim 21, Blumenau fails to disclose sending a lock request to a master port of a network device. Rather, Blumenau appears to disclose that locking and unlocking is performed by the cached storage subsystem, rather than contacting a master port of a network device.

The Examiner cites Brewer, asserting that Brewer teaches a method, wherein requesting a lock of the one or more physical storage locations comprises: sending a lock request to a master port of a network device within the storage area network, wherein the master port is adapted for managing lock requests. The Examiner cites the abstract and col. 8, lines 20-60 of Brewer. However, Applicant was unable to find any reference to requesting a lock or sending a lock request to a master port. Rather, it appears that these portions of Brewer disclose redirection rather than sending a lock request (e.g., to prevent other ports or network devices from accessing or altering the locked storage locations). Redirection is an entirely different concept from locking. Moreover, Brewer fails to cure the deficiencies of Blumenau and Latif. Accordingly, Applicant respectfully submits that claims 21-28 are patentable over the cited references.

Applicant respectfully submits that the independent claims are patentable over the cited references, separately or in combination. The dependent claims depend from one of the independent claims and are therefore patentable for at least the same reasons. However, the dependent claims recite additional limitations that further distinguish them from each of the cited references. The additional limitations recited in the independent claims or the dependent claims are not further discussed, as the above discussed limitations are clearly sufficient to distinguish the claimed invention from the cited reference. Thus, it is respectfully requested that the Examiner withdraw the rejection of the claims under 35 USC §103(a).

RECEIVED
CENTRAL FAX CENTER

NOV 17 2006

SUMMARY

If there are any issues remaining which the Examiner believes could be resolved through either a Supplemental Response or an Examiner's Amendment, the Examiner is respectfully requested to contact the undersigned attorney at the telephone number listed below.

Applicants hereby petition for an extension of time which may be required to maintain the pendency of this case, and any required fee for such extension or any further fee required in connection with the filing of this Amendment is to be charged to Deposit Account No. 50-0388 (Order No. ANDIP003).

Respectfully submitted,
BEYER, WEAVER & THOMAS, LLP


Elise R. Heilbrunn
Reg. No. 42,649

BEYER, WEAVER & THOMAS, LLP
P.O. Box 70250
Oakland, CA 94612-0250

Tel: (510) 663-1100